

# Introduction to CloudXPRT

CloudXPRT is a cloud benchmark that can accurately measure the performance of modern, cloud-first applications deployed on modern infrastructure as a service (IaaS) platforms, whether those platforms are paired with on-premises (datacenter), private cloud, or public cloud deployments.

July 2, 2020

# CloudXPRT

**BenchmarkXPRT**

BenchmarkXPRT Development Community

# Table of contents

- Introduction .....3
- The workloads .....3
  - Web microservices.....3
  - Data analytics .....4
- Choosing and downloading CloudXPRT installation packages.....4
  - Installation packages.....4
  - Minimum system requirements .....4
  - Where to find CloudXPRT .....4
- Running CloudXPRT.....5
  - Web microservices test instructions and information .....5
  - Data analytics test instructions and information.....5
- After running CloudXPRT.....5
  - Test results files .....5
  - Submitting test results .....5
  - The results review group and submission schedule .....6
    - The results review group .....6
    - The submission, review, and publication cycle.....7
    - The publication schedule .....7
  - Viewing results from other testers .....7
- Accessing the source code .....7
- Future development.....8
  - New workloads .....8
  - Future software and installer package updates.....8
- Conclusion .....8
- About the BenchmarkXPRT family.....9
  - The community model .....9

# Introduction

This paper provides an introductory overview of CloudXPRT. CloudXPRT is a cloud benchmark that can accurately measure the performance of applications deployed on modern infrastructure as a service (IaaS) platforms, whether those platforms are paired with on-premises (datacenter), private cloud, or public cloud deployments. Applications increasingly use clouds in latency-critical, highly available, and high-compute scenarios, so we designed CloudXPRT to use cloud-native components on an actual stack to produce end-to-end performance metrics that can help users determine the right IaaS configuration for their business.

## CloudXPRT

- Is compatible with on-premises, private, and public cloud deployments
- Runs on top of cloud platform software such as [Kubernetes](#) and [Docker](#)
- Supports multi-tier workloads
- Reports relevant metrics such as throughput and critical latency for responsiveness-driven applications, and maximum throughput for applications dependent on batch processing

In this document, you will find information about CloudXPRT workloads, choosing and downloading installation packages, and submitting CloudXPRT results for publication. In future whitepapers, we will discuss in greater detail each of the benchmark workloads, the range of test configuration options, and results reporting and methods for analysis.

## The workloads

CloudXPRT currently supports two workloads: Web microservices and Data analytics. Testers can use metrics from the workloads to compare IaaS stack (both hardware and software) performance and to evaluate whether any given stack is capable of meeting service-level agreement (SLA) thresholds. Testers can run CloudXPRT on local datacenter, Amazon Web Services (AWS), Google Cloud Platform (GCP), or Microsoft Azure deployments.

### **WEB MICROSERVICES**

In the Web microservices workload, a simulated user logs into a web application that does three things: provides a selection of stock options, performs Monte-Carlo simulations with those stocks, and presents the user with options that may be of interest. This scenario enables the workload to model a traditional three-tier web application with services in the web, application, and data layers. The workload uses Kubernetes, Docker, NGINX, REDIS, Cassandra, and monitoring modules to mimic an end-to-end IaaS scenario.

The workload reports performance in transactions per second, which reflects the number of successful requests per second the stack achieves for each level of concurrency. Testers can use this workload's metrics to compare IaaS stack performance and to evaluate whether any given stack is capable of meeting SLA thresholds.

## DATA ANALYTICS

The Data analytics workload calculates XGBoost model training time. XGBoost is a gradient-boosting framework that data scientists often use for machine learning (ML)-based regression and classification problems. The purpose of the workload in the context of CloudXPRT is to measure how well an IaaS stack enables XGBoost to speed and optimize model training. To do this, the data analytics workload uses Kubernetes, Docker, object storage, message pipeline, and monitorization modules to mimic an end-to-end IaaS scenario.

The workload reports latency in the form of job duration in seconds, and throughput rate in the form of jobs per minute. Testers can use this workload's metrics to compare IaaS stack performance and to evaluate whether any given stack is capable of meeting SLA thresholds.

## Choosing and downloading CloudXPRT installation packages

### INSTALLATION PACKAGES

To simplify installation, we separated the CloudXPRT installation packages by workload and target platform. There are five installation packages: a single Data analytics installation package, which is compatible with all four target platforms (on-premises, AWS, Azure, and GCP), and four Web microservices installation packages, one for each of the four target platforms.

### MINIMUM SYSTEM REQUIREMENTS

We designed CloudXPRT for high-end servers. During each workload, the benchmark scales to utilize all the available cores. However, for functional testing, the physical node(s) or VM(s) under test must have at least:

- Ubuntu 18.04 LTS
- 16 logical or virtual CPUs
- 8 GB RAM
- 10 GB of available disk space (20 GB for the data analytics workload)
- An internet connection

For three of the four target platforms, on-premises, AWS, and GCP, testing requires both Docker and Kubernetes. Testing on Microsoft Azure requires only Kubernetes.

Off-premises tests will require access to an AWS, Azure, or GCP account, depending on the desired test configuration.

### WHERE TO FIND CLOUDXPRT

All of the CloudXPRT installation packages are available through the CloudXPRT [download table](#) (see screenshot below) and through the CloudXPRT GitHub [repository](#).

Workload	Target platform	Install package	Documentation
Data analytics	On premises/AWS/Azure/GCP	<a href="#">Download</a>	<a href="#">Readme</a>
Web microservices	On premises	<a href="#">Download</a>	<a href="#">Readme</a>
Web microservices	Amazon Web Services (AWS)	<a href="#">Download</a>	<a href="#">Readme</a>
Web microservices	Google Cloud Platform (GCP)	<a href="#">Download</a>	<a href="#">Readme</a>
Web microservices	Microsoft Azure	<a href="#">Download</a>	<a href="#">Readme</a>

## Running CloudXPRT

We present detailed installation, setup, and test configuration instructions for each of CloudXPRT's workloads in separate readme documents located in the [CloudXPRT GitHub repository](#).

### WEB MICROSERVICES TEST INSTRUCTIONS AND INFORMATION

[Click here](#) to access instructions for the CloudXPRT web microservices workload.

### DATA ANALYTICS TEST INSTRUCTIONS AND INFORMATION

[Click here](#) to access instructions for the CloudXPRT data analytics workload.

## After running CloudXPRT

### TEST RESULTS FILES

CloudXPRT produces local copies of results summary, log, and input run configuration files at the end of each test run. Testers must run a parsing script to generate the results summary file for the data analytics workload, but the benchmark automatically generates the other files.

### SUBMITTING TEST RESULTS

We invite and encourage everyone to submit benchmark results from their testing for inclusion in the public CloudXPRT results table. To do so, please follow the steps below.

1. Create a folder to collect the files you will submit.
2. After a benchmark run completes, copy the following items to the folder you created in step 1:  
For the Web microservices workload
  - The CSV results file located at CloudXPRT\_vXXX\_web-microservices/cnbrun/output/autoloader\_mc\_YYYYMMDD\_\*.csv.
  - The log files located at CloudXPRT\_vXXX\_web-microservices/cnbrun/output/autoloader\_mc\_YYYYMMDD\_\*.log.
  - The input run configuration file located at CloudXPRT\_vXXX\_web-microservices/cnbrun/config.json.

For the Data analytics workload

- The CSV results file generated at CloudXPRT\_vXXX\_data-analytics/cnbrun/results.csv, which was generated by the command line `./cnb-analytics_parse-all-results.sh | sed -e 's/\s\+/,/g' > results.csv`.
  - The log files generated at CloudXPRT\_vXXX\_data-analytics/cnbrun/output/autoloader\_xgboost\_all\_YYYYMMDD\_\*.log.
  - The input run configuration file located at CloudXPRT\_vXXX\_data-analytics/cnbrun/cnb-analytics\_config.json.
3. Locate the SystemInfo.csv file in the root directory of the installation package you used, and open it.
  4. Complete the required fields, and describe any changes you made to customize scripts, models, or files for result generation. (Note: We collect this information to make it possible for others to reproduce the test and confirm that they get similar results.)
  5. Save the SystemInfo.csv file to the folder you created in step 1, and create a single zip file containing all of the files in this folder.
  6. Create an email message to the BenchmarkXPRT Community Administrator:
    - Send to [benchmarkxpertsupport@principledtechnologies.com](mailto:benchmarkxpertsupport@principledtechnologies.com).
    - Confirm that the reply-to address you specify is a valid address inside your organization.
    - Use the subject “CloudXPRT Results Submission.”
    - In the body of the message, specify your company name and name of the person responsible for the test.
    - Attach the zip file you created in step 5.

When we receive your submission, we will verify the tester’s identity and validate the results. Due to the complexity of CloudXPRT tests, and to be as transparent and accurate as possible with our published results, we may ask follow-up questions about the tests and/or system configuration. If we decide to include your results in the public database, we will notify you.

## THE RESULTS REVIEW GROUP AND SUBMISSION SCHEDULE

We are using a periodic results submission process for CloudXPRT. Each month, a results review group examines results submissions, and we publish all approved results on a scheduled day.

While testers are free to publish results outside of the monthly review process, we only publish results on CloudXPRT.com after first vetting them through the formal process. Our goal is to strike a balance between allowing the tech press, vendors, or other testers to publish CloudXPRT results on their own schedule, and simultaneously building a curated results database that OEMs or other parties can use to compete for the best results.

### The results review group

The CloudXPRT results review group serves as a sanity check and a forum for comments on each month’s submissions. All registered BenchmarkXPRT Development Community members who wish to participate in the review process can join the group by contacting us via [email](#). We’ll confirm receipt of your request and add you to the review group mailing list. Any non-members who would like to join the review group can contact us and we’ll help you become community members.

### The submission, review, and publication cycle

We update the CloudXPRT results database once a month on a published schedule. While testers can submit results through the CloudXPRT [results submission page](#) at any time, two weeks prior to each publication date, we close submissions for that review cycle. One week prior to each publication date, we email details of that month's submissions to the results review group, along with the deadline for sending post-publication feedback.

### The publication schedule

We publish results to the database on the last business day of each month and close the submission window at 11:59 PM on the business day that falls two weeks earlier (with occasional adjustments for holidays). The schedule is available at least six months in advance on CloudXPRT.com. Due to holiday schedules, we do not publish results in December.

## VIEWING RESULTS FROM OTHER TESTERS

At CloudXPRT.com, interested parties can view [test results](#) published by the BenchmarkXPRT Development Community. The following tips will help visitors navigate the results viewer:

- Click the tabs at the top of the table to switch from Data analytics workload results to Web microservices workload results.
- Click the header of any column to sort the data on that variable. Single click to sort A to Z and double-click to sort Z to A.
- Click the link in the Source/details column to visit a detailed page for that result, where you'll find additional test configuration and system hardware information and the option to download results files.
- By default, the viewer displays eight results per page, which you can change to 16, 48, or Show all.
- The free-form search field above the table lets you filter for variables such as cloud service or processor.

We will add more features, including expanded filtering and sorting mechanisms, to the results viewer in the near future. We are also investigating ways to present multiple data points in a graph format, which will allow visitors to examine performance behavior curves in conjunction with factors such as concurrency and resource utilization.

## Accessing the source code

We expect to make the CloudXPRT source code available to the public in July 2020 via the [CloudXPRT GitHub repository](#). As we've discussed in the past, publishing XPRT source code is part of our commitment to making the XPRT development process as transparent as possible. With some XPRT benchmarks, we've made the source code available to only community members. With CloudXPRT, we have released the source code more widely. By allowing all interested parties, not just community members, to download and review our source code, we're taking tangible steps to improve openness and honesty in the benchmarking industry and are encouraging the kind of constructive feedback that helps to ensure that the XPRTs continue to contribute to a level playing field. Traditional open-source models encourage developers to change products and even take

them in new and different directions. Because benchmarking requires a product that remains static to enable valid comparisons over time, we allow people to download the source code and submit potential workloads for future consideration, but we reserve the right to control derivative works. This discourages a situation where someone publishes an unauthorized version of the benchmark and calls it an “XPRT.” We encourage you to download and review the source and send us any feedback you may have. Your questions and suggestions may influence future versions of CloudXPRT.

## Future development

### NEW WORKLOADS

We hope to add at least one new workload to the CloudXPRT suite in the coming year, and are especially interested in the possibility of adding a container scaling workload. We’re always open to new workload ideas, so if you have a suggestion, please contact us at [BenchmarkXPRTsupport@principledtechnologies.com](mailto:BenchmarkXPRTsupport@principledtechnologies.com).

### FUTURE SOFTWARE AND INSTALLER PACKAGE UPDATES

With several cloud management components, two workloads, three public cloud service providers, and a wide range of possible configurations and use cases, CloudXPRT has more moving parts than most of the legacy XPRT benchmarks. For this reason, and because we want CloudXPRT to provide consistently relevant evaluation data in the complex cloud computing space, we anticipate updating CloudXPRT more frequently than we have other XPRTs. With that in mind, we want to assure CloudXPRT testers that when we update CloudXPRT, they can expect minimized disruption, consideration for their testing needs, and clear communication.

Each CloudXPRT component application (Docker, Kubernetes, etc.) has its own development schedule, and new versions can appear with minimal advance notice. In addition, at any time, the three CloudXPRT-compatible cloud service providers can make changes that affect CloudXPRT installation, configuration, testing, and results. Because the CloudXPRT installation packages are modular, we’ll be able to update individual packages when necessary. If you test with only a single CloudXPRT package, updates to the other packages won’t affect your testing.

It’s not possible to both maintain CloudXPRT’s relevance and avoid all disruption, but we’ll work to keep it to a minimum. When software compatibility issues force us to update a CloudXPRT package, we will test to see whether results remain comparable to those from previous versions. As always, we will use documentation and versioning to reduce confusion and keep testers informed. We’ll document all changes thoroughly in the package readmes, and will discuss significant updates in the XPRT blog

## Conclusion

We hope this paper has answered any initial questions you may have about CloudXPRT. For more information, visit us at [CloudXPRT.com](http://CloudXPRT.com) and [BenchmarkXPRT.com](http://BenchmarkXPRT.com). If you cannot find the answer to your question, you need help with CloudXPRT, or you have suggestions on ways to improve CloudXPRT, send an email to our team at [BenchmarkXPRTsupport@principledtechnologies.com](mailto:BenchmarkXPRTsupport@principledtechnologies.com).



## About the BenchmarkXPRT family

The BenchmarkXPRT tools are a set of apps that help you test how well devices do the kinds of things you do every day. In addition to CloudXPRT, the BenchmarkXPRT suite currently comprises the following tools:

- [AIXPRT](#), an AI benchmark tool that makes it easier to evaluate a system's machine learning inference performance by running common image-classification, object-detection, and recommender system workloads.
- [WebXPRT](#), a browser benchmark that compares the performance of almost any web-enabled device
- [CrXPRT](#), an app to test the responsiveness and battery life of Chromebooks
- [HDXPRT](#), a benchmark to test how well Windows PCs handle real-world apps
- [TouchXPRT](#), a Universal Windows Platform app to test the responsiveness of Windows 10 devices
- [MobileXPRT](#), an app to test the responsiveness of Android devices

We designed the apps to test a wide range of devices on a level playing field. When you look at results from XPRTs, you get unbiased, fair product comparison information.

### THE COMMUNITY MODEL

We built BenchmarkXPRT around a unique community model. Community membership is open to anyone, and there are many ways to participate. Members of the BenchmarkXPRT Development Community are involved in every step of the process. They give input on the design of upcoming versions, contribute source code, and help test the resulting implementation. Community members also have access to early releases in the form of community previews.

The community helps us avoid the ivory tower syndrome. Diversity of input during the design process makes the tests more representative of real-world activity. Giving community members access to the source code both improves the implementation of the design and increases confidence in the code.

The community model differs from the open source model primarily by controlling derivative works. It is important that the BenchmarkXPRT benchmarks return consistent results. If the testing community calls different derivative works by the same name, the result would be that test results would not be comparable. That would limit, if not destroy, the tools' effectiveness.

If you are not currently a community member, we encourage you to join! Our community is open to everyone, from software developers to interested consumers. Not only will you get early releases of future XPRTs, but you will also be able to influence the future of the tools. [Register](#) now, or for more information, see the [BenchmarkXPRT FAQ](#).